arXiv:1201.5991v1 [math.NA] 28 Jan 2012

Parallel implementation of Multilevel BDDC

Jakub Šístek, Jan Mandel, Bedřich Sousedík, Pavel Burda

Abstract In application of the Balancing Domain Decomposition by Constraints (BDDC) to a case with many substructures, solving the coarse problem exactly becomes the bottleneck which spoils scalability of the solver. However, it is straightforward for BDDC to substitute the exact solution of the coarse problem by another step of BDDC method with subdomains playing the role of elements. In this way, the algorithm of three-level BDDC method is obtained. If this approach is applied recursively, multilevel BDDC method is derived. We present a detailed description of a recently developed parallel implementation of this algorithm. The implementation is applied to an engineering problem of linear elasticity and a benchmark problem of Stokes flow in a cavity. Results by the multilevel approach are compared to those by the standard (two-level) BDDC method.

Jakub Šístek
Institute of Mathematics, Academy of Sciences of the Czech Republic,
Žitná 25, Praha 1, CZ-115 67, Czech Republic, e-mail: sistek@math.cas.cz

Jan Mandel
Department of Mathematical and Statistical Sciences, University of Colorado Denver,
Campus Box 170, Denver, CO 80217-3364, United States, e-mail: jan.mandel@ucdenver.edu

Bedřich Sousedík
Department of Aerospace and Mechanical Engineering, University of Southern California,
Olin Hall 430, Los Angeles, CA 90089-2531, United States, e-mail: sousedik@usc.edu

Pavel Burda
Department of Mathematics, Faculty of Mechanical Engineering, Czech Technical University,
Karlovo náměstí 13, Praha 2, CZ-121 35, Czech Republic, e-mail: pavel.burda@fs.cvut.cz

1 Introduction

The Balancing Domain Decomposition by Constraints (BDDC) method introduced in [2] is one of the most advanced methods of iterative substructuring for the solution of large systems of linear algebraic equations arising from discretization of boundary value problems. However, in the case of many substructures, solving the coarse problem exactly becomes the limiting factor for scalability of the otherwise perfectly parallel algorithm. This has been observed also for the FETI-DP method (e.g. in [3]), which is closely related to BDDC. For this reason, recent research in the area is directed towards inexact solutions of the coarse problem. For example, algebraic multigrid is used in [3] to obtain an approximate coarse correction within the FETI-DP method, and excellent scalability is achieved.

We follow a different approach in this contribution. As was mentioned already in [2], it is quite straightforward for BDDC to substitute the exact solution of the coarse problem by another step of the BDDC method with subdomains playing the role of elements. In this way, the algorithm of three-level BDDC method is obtained (studied in [8]). If this step is repeated recursively, one arrives at the *multilevel BDDC method* (introduced in [4] without a parallel implementation). Unlike for most other domain decomposition methods, such extension is natural for BDDC, since the coarse problem has the same structure as the original problem. Although the mathematical theory in [4] suggests worsening of the efficiency of the multilevel BDDC preconditioner with each additional level, the resulting algorithm may outperform the standard method with respect to computational time due to better scalability. This fact makes the algorithm a good candidate for using on future massively parallel systems.

In this paper, we present a recently developed parallel implementation of multilevel BDDC method. It is applied to an engineering problem of linear elasticity and a benchmark problem of Stokes flow in a cavity. The results suggest which drawbacks of the two-level implementation might be overcome by the extension to more levels. Our solver library has been released as an open-source package.

2 BDDC preconditioner with two and more levels

The starting point for BDDC is the *reduced interface problem* $\widehat{\mathbf{S}}\widehat{\mathbf{u}} = \widehat{\mathbf{g}}$, where $\widehat{\mathbf{S}}$ is the *Schur complement* with respect to *interface*, i.e. unknowns shared by more than one subdomain, $\widehat{\mathbf{u}}$ is the part of vector of coefficients of finite element basis functions at the interface, and $\widehat{\mathbf{g}}$ is sometimes called *condensed right hand side*. This problem is solved by a Krylov subspace method in the framework of *iterative substructuring*. Within these methods, application of $\widehat{\mathbf{S}}$ to a vector is realized by parallel solution of independent *discrete Dirichlet problems*. In this way, the costly explicit construction of the Schur complement is avoided. However, since it is not the main concern of this contribution, the reader is referred to paper [5], or monograph [7] for details of iterative substructuring.

In what follows, we turn our attention towards the second key part of Krylov subspace methods – the *preconditioner*, which is realized by one step of the BDDC method. Let us begin with description of the standard (two-level) version of BDDC. Let $\mathbf{K}_i$ be the local subdomain matrix, obtained by the sub-assembling of element matrices of elements contained in i-th subdomain. We introduce the *coarse space basis functions* on each subdomain represented by columns of matrix Ψ_i, which is the solution to the saddle point problem with multiple right hand sides

$$\begin{bmatrix} \mathbf{K}_i & \mathbf{C}_i^T \\ \mathbf{C}_i & \mathbf{0} \end{bmatrix} \begin{bmatrix} \Psi_i \\ \Lambda_i \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{I} \end{bmatrix}. \tag{1}$$

Matrix $\mathbf{C}_i$ represents constraints on functions Ψ_i, one row per each. These constraints enforce continuity of approximate solution at *corners* and/or continuity of more general quantities, such as averages over shared subsets of interface (*edges* or *faces*) between adjacent subdomains. The *local coarse matrix* $\mathbf{K}_{Ci} = \Psi_i^T \mathbf{K}_i \Psi_i = -\Lambda_i$ is constructed for each subdomain. The *global coarse matrix* $\mathbf{K}_C$ is obtained by the assembly procedure from local coarse matrices. This can be formally written as $\mathbf{K}_C = \sum_{i=1}^N \mathbf{R}_{Ci}^T \mathbf{K}_{Ci} \mathbf{R}_{Ci}$, where $\mathbf{R}_{Ci}$ realize the restriction of global coarse degrees of freedom to local coarse degrees of freedom of i-th subdomain.

Suppose $\widehat{\mathbf{r}} = \widehat{\mathbf{g}} - \widehat{\mathbf{S}}\widehat{\mathbf{u}}$ is a residual within the Krylov subspace method. The residual assigned to i-th subdomain is computed as $\mathbf{r}_i = \mathbf{E}_i^T \widehat{\mathbf{r}}$, where matrices of weights $\mathbf{E}_i^T$ distribute $\widehat{\mathbf{r}}$ to subdomains. The subdomain correction is now defined as the solution to the system

$$\begin{bmatrix} \mathbf{K}_i & \mathbf{C}_i^T \\ \mathbf{C}_i & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{z}_i \\ \lambda_i \end{bmatrix} = \begin{bmatrix} \mathbf{r}_i \\ \mathbf{0} \end{bmatrix}. \tag{2}$$

The residual for the coarse problem is constructed using the coarse basis functions subdomain by subdomain and assembling the contributions as $\mathbf{r}_C = \sum_{i=1}^N \mathbf{R}_{Ci}^T \Psi_i^T \mathbf{E}_i^T \widehat{\mathbf{r}}$. The coarse correction is defined as the solution to problem $\mathbf{K}_C \mathbf{z}_C = \mathbf{r}_C$. Both corrections are finally added together and averaged on the interface by matrices $\mathbf{E}_i$ to produce the preconditioned residual $\widehat{\mathbf{z}} = \sum_{i=1}^N \mathbf{E}_i (\Psi_i \mathbf{R}_{Ci} \mathbf{z}_C + \mathbf{z}_i)$.

In the three-level BDDC method [8], the matrix $\mathbf{K}_C$ is not constructed on the second level. Instead, subdomains from the basic (first) level are grouped into subdomains on the next (second) level in the same way as elements of the original mesh are grouped into subdomains of the first level. The whole procedure described in this section is now repeated for the second level and thus the final coarse problem represents the third level. Obviously, this can be repeated again in the multilevel BDDC method. The only important difference between the first and the higher levels is the additional *interior pre-correction* and *post-correction* applied on higher levels in order to approximate the whole vector of coarse solution on the lower level.

According to [4], the condition number of the operator preconditioned by multilevel BDDC with L levels satisfies $\kappa(\mathbf{M}_{BDDC}\widehat{\mathbf{S}}) \leq \prod_{\ell=1}^{L-1} C_\ell \left(1 + \log \frac{H_\ell}{H_{\ell-1}}\right)^2$, where H_ℓ is the characteristic size of subdomain on level ℓ, and $H_0 \equiv h$ is the characteristic size of element. Index ℓ is used here and throughout the next section to denote particular level. Due to the product present in this bound, each additional level worsens

the mathematical efficiency of the multilevel preconditioner. The proof of the condition number bound as well as details of the algorithm of multilevel BDDC can be found in [4].

3 Parallel implementation

Our implementation of the multilevel BDDC method has been recently released as an open-source solver library BDDCML[1]. It is written in Fortran 95 programming language and parallelized by MPI. The solver relies on the sparse direct solver MUMPS — a serial instance is used for each subdomain problem and a parallel instance is called for the final coarse problem. The solver supports assignment of several subdomains to each processor, since it is often useful to create divisions independently of number of available processors. A division of the mesh into subdomains on the first level is either provided to the solver by user's application or created internally by ParMETIS. The METIS package is currently used for this purpose on higher levels.

Similarly to other related preconditioners, we first need to set-up the multilevel BDDC preconditioner, which is then applied in each iteration of the Krylov subspace method. Details of the set-up are given in Algorithm 1, while key operations of each application are summarized in Algorithm 2. In these descriptions, we provide comments on how the steps are implemented in BDDCML in parentheses.

Algorithm 1 Set-up of BDDC preconditioner with L levels

for level $\ell = 1, \ldots, L-1$ **do**
  if $\ell > 1$ **then**
    build *pseudo-mesh*: subdomains $\rightarrow$ 'elements'; corners + edges + faces $\rightarrow$ 'nodes'
  end if
  divide *pseudo-mesh* into subdomains (by METIS for $\ell > 1$, or by ParMETIS for $\ell = 1$)
  classify interface into *faces*, *edges*, *vertices*
  select *corners* (using face-based algorithm from [6])
  assemble matrices of subdomains $\mathbf{K}_i^\ell$ (use MPI to collect them on assigned cores)
  prepare interior correction – factorize interior block of $\mathbf{K}_i^\ell$ (serial MUMPS)
  factorize the matrices of local saddle point problems (1) (serial MUMPS)
  find *coarse basis functions* Ψ_i^ℓ and *coarse matrices* $\mathbf{K}_{Ci}^\ell = -\Lambda_i^\ell$ from (1) (serial MUMPS)
end for
factorize *global coarse matrix* $\mathbf{K}_C^{L-1} = \sum_{i=1}^{N_{L-1}} (\mathbf{R}_{Ci}^{L-1})^T \mathbf{K}_{Ci}^{L-1} \mathbf{R}_{Ci}^{L-1}$ (parallel MUMPS)

[1] `http://www.math.cas.cz/$\sim$sistek/software/bddcml.html`

Algorithm 2 Application of BDDC preconditioner with L levels

for level $\ell = 1,\ldots,L-1$ **do**
  if $\ell > 1$ **then**
    $\widehat{\mathbf{r}}^{\ell} \leftarrow \mathbf{r}_C^{\ell-1}$
    compute *interior pre-correction* of residual $\widehat{\mathbf{r}}^{\ell}$ (serial MUMPS)
  end if
  distribute residual among subdomains $\mathbf{r}_i^{\ell} = (\mathbf{E}_i^{\ell})^T \widehat{\mathbf{r}}^{\ell}$
  determine subdomain corrections $\mathbf{z}_i^{\ell}$ from (2) (serial MUMPS)
  construct coarse residual $\mathbf{r}_C^{\ell} = \sum_{i=1}^{N_\ell} (\mathbf{R}_{Ci}^{\ell})^T (\Psi_i^{\ell})^T (\mathbf{E}_i^{\ell})^T \widehat{\mathbf{r}}^{\ell}$ (collective MPI)
end for
solve the coarse problem $\mathbf{K}_C^{L-1}\,\mathbf{z}_C^{L-1} = \mathbf{r}_C^{L-1}$ (parallel MUMPS)
for level $\ell = L-1,\ldots,1$ **do**
  if $\ell < L-1$ **then**
    $\mathbf{z}_C^{\ell} \leftarrow \widehat{\mathbf{z}}^{\ell+1}$
  end if
  combine coarse correction and subdomain corrections $\widehat{\mathbf{z}}^{\ell} = \sum_{i=1}^{N_\ell} \mathbf{E}_i^{\ell} \left(\Psi_i^{\ell} \mathbf{R}_{Ci}^{\ell} \mathbf{z}_C^{\ell} + \mathbf{z}_i^{\ell}\right)$
  if $\ell > 1$ **then**
    apply *interior post-correction* to $\widehat{\mathbf{z}}^{\ell}$ (serial MUMPS)
  end if
end for

4 Numerical results

The first example corresponds to a problem of mechanical analysis of a cubic sample of geocomposite and was analyzed in [1]. The length of the edge of the cube is 75 mm. The cube comprises five distinct materials identified by means of computer tomography (Fig. 1 left), which causes anisotropic response of the cube even for simple axial stretching in z direction (Fig. 1 right). The problem is discretized using unstructured grid of about 12 million linear tetrahedral elements, resulting in approximately 6 million unknowns. The mesh was divided into 1024, 128, and 16 subdomains on the first, second and third level, respectively, and the respective coarse problems (using corners and arithmetic averages on all edges and faces) contain 86,094, 11,265, and 612 unknowns.

Table 1 summarizes the efficiency of the multilevel preconditioner by means of the resulting condition number (estimated from the tridiagonal matrix generated during iterations of preconditioned conjugate gradient (PCG) method) and number of iterations. The iterations were stopped when the relative residual $\|\widehat{\mathbf{r}}\|/\|\widehat{\mathbf{g}}\|$ decreased bellow 10^{-6}. This table confirms the predicted worsening of the condition number with each additional level expected from the condition number bound.

Table 2 contains a strong scaling test using different number of levels. We differentiate the time spent on set-up and in PCG. All these computations were performed on the IBM SP6 computer at CINECA, Bologna. The computer is based on IBM Power6 4.7 GHz processors with 4 GB of RAM per core.

We can conclude from Tab. 2 that while adding levels seems not to be feasible for small number of cores (the computational time stagnates or even grows), it improves

the scaling on many cores. The minimal overall solution time is achieved for four levels and largest number of cores, despite the largest number of required iterations.

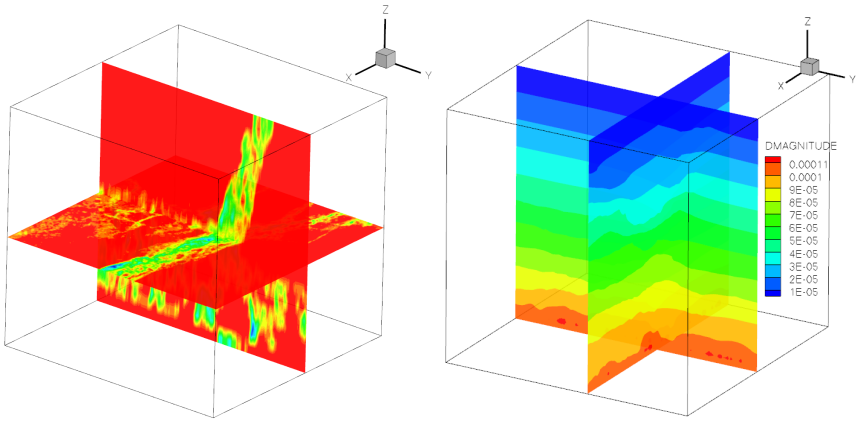

Fig. 1 Geocomposite problem: slices through material distribution (left) and displacement field (right).

Table 1 Condition number and number of iterations for different number of levels.

num. of levels	num. of subs.	cond. num.	num. of PCG its.
2	1024/1	50	46
3	1024/128/1	79	56
4	1024/128/16/1	568	131

Table 2 Strong scaling for geocomposite problem using two, three, and four levels.

number of processors	64	128	256	512	1024
2 levels					
BDDC set-up time (s)	61.0	37.7	25.7	23.2	39.5
PCG time (s)	22.3	19.9	27.8	44.9	97.5
3 levels					
BDDC set-up time (s)	49.5	29.0	18.4	12.6	11.0
PCG time (s)	28.5	22.6	16.7	14.7	13.2
4 levels					
BDDC set-up time (s)	49.4	28.6	17.8	12.3	9.1
PCG time (s)	60.6	33.2	21.2	15.4	11.8

Our second example is a problem of Stokes flow in a 3D lid driven cavity. We use the set-up suggested in [9]: zero velocity is prescribed on all faces of the $[0,1]^3$ cube except the face for z = 1, where unit velocity vector $\mathbf{u} = [1/\sqrt{3}, \sqrt{2/3}, 0]$ is

prescribed. We have used this test case also in the recent paper [5], but we have not presented parallel results there.

The problem is uniformly discretized using hexahedral Taylor–Hood finite elements. Computational mesh was divided into irregular partitions using the METIS graph partitioner (see Fig. 2 for an example). A plot of pressure inside the cavity and velocity vectors is given in Fig. 2 (right).

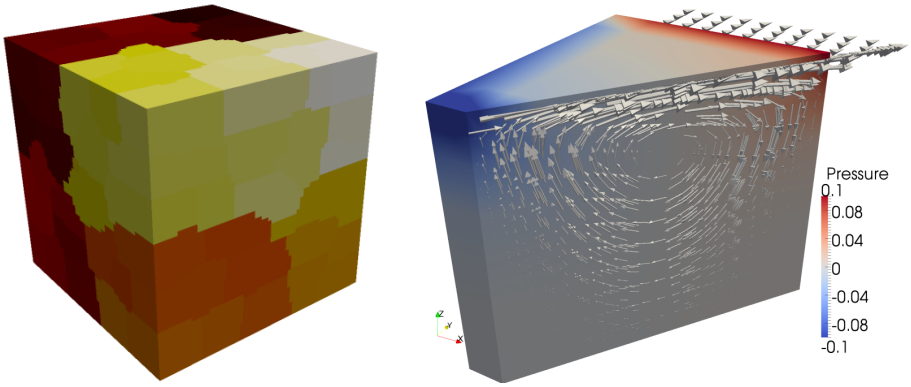

Fig. 2 Stokes flow in lid driven cavity: example of division into 64 subdomains (left), pressure contours and velocity vectors in the cut in direction of the prescribed velocity $[1/\sqrt{3}, \sqrt{2/3}, 0]$ (right).

Table 3 summarizes a weak scaling test for this problem. The sequence of problems ranging from 1.7 million unknowns to 25.4 million unknowns are distributed among processors such that the size of local problems is kept approximately constant around 50 thousand unknowns. We present results by the BDDC preconditioner using two and three levels combined with BiCGstab method. For comparison, we also report results of the preconditioner based on Additive Schwarz Method (ASM) combined with BiCGstab and using MUMPS offered in the PETSc library version 3.1. We report numbers of iterations and required overall computational times. Where 'n/a' is present in the table, the simple ASM method was unable to provide the solution, which was often encountered for larger problems. The computational times were obtained on Darwin supercomputer of the University of Cambridge, using Intel Xeon 5100 3.0 GHz processors with 2 GB of RAM per core.

We can see in Tab. 3 that number of BiCGstab iterations remains almost constant for the two-level method, while mildly growing when using three levels, being again larger for the latter. The computational time slightly grows with problem size for BDDC, both using two and three levels, but this growth is rather acceptable. More importantly, we can also see that for this case the benefit of using an additional level is slightly outweighed by the overhead of the additional iterations, and so the computational time is not improved by using three levels. Finally, the comparison with ASM results seem to confirm the computational efficiency of our solver

which is comparable with the state-of-the-art PETSc library, and the rather poor performance of ASM in comparison to BDDC in this setting.

Table 3 Weak scaling for Stokes flow in the cavity: additive Schwarz method (ASM), and BDDC using two and three levels.

			ASM		BDDC (2 levels)		BDDC (3 levels)		
# elms.	# dofs.	# cores	# its.	time (s)	# its.	time (s)	divisions	# its.	time (s)
40^3	1.7M	32	282	533	18	122	32/4/1	22	126
50^3	3.2M	64	396	805	19	132	64/8/1	25	205
64^3	6.7M	128	384	536	21	186	128/16/1	30	194
80^3	13.1M	256	n/a	n/a	21	178	256/32/1	36	201
100^3	25.4M	512	n/a	n/a	20	205	512/64/1	35	211

5 Conclusion

We have presented a parallel open-source implementation of the multilevel BDDC method. The two-level algorithm has scalability issues related to the coarse problem solution, mainly in the part of iterations. It can be noted that for the tested cases, it has not been the size of the coarse problem, but rather its fragmentation among too many cores which causes these issues. From our experiments, it appears that the multilevel preconditioner tends to scale better in both parts – set-up and Krylov subspace iterations. While the better scalability is able to translate into much faster solution for some cases, the extra overhead can also just cancel out the savings for other cases. It is therefore important to choose appropriate number of levels for a particular problem. We expect that advantages of the multilevel approach would pronounce further for problems divided into many (tens of thousands) of subdomains. Such challenging problems will likely become common in near future and will provide valuable feedback for further research in this field.

Acknowledgements We are grateful to Prof. Blaheta and Dr. Starý (Institute of Geonics AS CR) for providing the geocomposite problem. We are also grateful to Dr. Cirak (University of Cambridge) for providing the OpenFTL package and computer time on Darwin. This work was supported by Ministry of Education, Youth and Sports of the Czech Republic under research project LH11004, by Czech Science Foundation under project 106/08/0403, by Institutional Research Plan AV0Z 10190503 of the AS CR, by grant IAA100760702 of the Grant Agency of AS CR, and by National Science Foundation under grant DMS-0713876. The research was started during two visits of Jakub Šístek at the University of Colorado Denver and some parts of the work have been performed under the HPC-Europa2 project with the support of the European Commission.

References

1. Blaheta, R., Jakl, O., Starý, J., Krečmer, K.: The Schwarz domain decomposition method for analysis of geocomposites. In: B. Topping et al. (eds.) Proceedings of the Twelfth International Conference on Civil, Structural and Environmental Engineering Computing. Civil-Comp Press, Stirlingshire, Scotland (2009)
2. Dohrmann, C.R.: A preconditioner for substructuring based on constrained energy minimization. SIAM J. Sci. Comput. **25**(1), 246–258 (2003)
3. Klawonn, A., Rheinbach, O.: Highly scalable parallel domain decomposition methods with an application to biomechanics. ZAMM Z. Angew. Math. Mech. **90**(1), 5–32 (2010)
4. Mandel, J., Sousedík, B., Dohrmann, C.R.: Multispace and multilevel BDDC. Computing **83**(2-3), 55–85 (2008)
5. Šístek, J., Sousedík, B., Burda, P., Mandel, J., Novotný, J.: Application of the parallel BDDC preconditioner to the Stokes flow. Comput. & Fluids **46**, 429–435 (2011)
6. Šístek, J., Čertíková, M., Burda, P., Novotný, J.: Face-based selection of corners in 3D substructuring. Math. Comput. Simulation DOI 10.1016/j.matcom.2011.06.007. Available online
7. Toselli, A., Widlund, O.B.: Domain Decomposition Methods—Algorithms and Theory, *Springer Series in Computational Mathematics*, vol. 34. Springer-Verlag, Berlin (2005)
8. Tu, X.: Three-level BDDC in three dimensions. SIAM J. Sci. Comput. **29**(4), 1759–1780 (2007)
9. Wathen, A.J., Loghin, D., Kay, D.A., Elman, H.C., Silvester, D.J.: A new preconditioner for the Oseen equations. In: F. Brezzi et al. (eds.) Numerical mathematics and advanced applications, pp. 979–988. Springer-Verlag Italia, Milano (2003). Proceedings of ENUMATH 2001